

REMARKS

Claims 1-120 are now pending in the application. The Examiner is respectfully requested to reconsider and withdraw the rejections in view of the amendments and remarks contained herein.

SPECIFICATION

The specification stands objected to as failing to provide proper antecedent basis for the claimed subject matter. Claims 35, 70, and 112 recite the phrase “first and second pages to be overwritten.” The Examiner alleges that the first and second pages are not discussed in the specification, citing MPEP § 608.01(o). Applicants respectfully disagree.

The cited portion of the MPEP states that the “meaning of every term used in any of the claims should be apparent from the descriptive portion of the specification.” Applicants note that the terms “first” and “second” to describe first and second pages are commonly used in claim language to distinguish separate elements. Here, Applicants are merely using the terms to distinguish separate pages. Support for separate pages can be found throughout Applicants’ drawings and specification. For example, Paragraph [0072] states “the least used page algorithm may be used to identify a least used page and a second least used page when a miss occurs. There is no requirement that the exact term “first and second pages” appears in the specification in MPEP § 608.01(o) or any other section.

CLAIMS OBJECTIONS

Claim 69 is objected to for certain informalities. Applicants amended claim 69 according to the Examiner’s suggestion. This amendment is not a narrowing amendment.

REJECTION UNDER 35 U.S.C. § 112

Claims 22, 32-34, 40, 42-43, 67-69, 75, 77-78, 109-111, and 117-120 are rejected under 35 U.S.C. § 112, first paragraph, as failing to comply with the enablement requirement. This rejection is respectfully traversed.

With respect to claims 32-34, 40, 42-43, 67-69, 75, 77-78, 109-111, and 117-120, the Examiner alleges that it is not clear “how one having ordinary skill in the art would be able to identify and replace pages requested by the second CPU by monitoring internal state of the first CPU.” Applicants respectfully note that the claims do not recite this alleged limitation.

For example, claim 32 recites receiving read requests from first and second CPUs and resolving line cache access conflicts between the first and second CPUs. Subsequently, claim 32 recites “a least used page device that identifies a least used page in said line cache, wherein said line cache module monitors at least one internal state of the first CPU, wherein said least used page device identifies a second least used page, and wherein said line cache state module selectively overwrites one of said first and second least used pages said based on internal state of the first CPU.” In other words, claim 32 recites identifying least used and second least used pages based on an internal state of the first CPU, but does not recite “requested by the second CPU” as the Examiner alleges. Applicants respectfully note that limitations that are not recited in the claims do not need to be enabled. As such, claims 32-34, 40, 42-43, 67-69, 75, 77-78, 109-111, and 117-120 are enabled.

With respect to claim 22, the Examiner alleges that “it is not clear how cache receives first and second addresses from arbitration device.” Applicants respectfully disagree. Claim 21 recites a line cache arbitration device that communicates with said first and second line cache interfaces and said line cache. Claim 22, which depends from claim 21, recites that the line cache receives said selected one of said first and second addresses from said line cache arbitration device. As shown in an exemplary embodiment in FIGS. 2 and 3, a line cache arbitration device 100 is in communication with a line cache 58. As such, Applicants respectfully submit that claim 22 is enabled.

Claims 6, 12, 15-20, 40, 51, 97, and 117 are rejected under 35 U.S.C. § 112, second paragraph, as being indefinite for failing to particularly point and distinctly claim the subject matter which Applicant regards as the invention. These rejections are respectfully traversed.

With respect to claim 6, the Examiner alleges that there is insufficient antecedent basis for the term “a second address.” Applicants respectfully note that antecedent basis is not required for a claim limitation that is newly introduced with an indefinite article (e.g. “a” or “an”). Applicants respectfully submit that claim 6 is definite.

With respect to claims 12, 15-19, 40, 51, and 117, Applicants amended the claims to include sufficient antecedent basis. These amendments are not narrowing amendments. Applicants respectfully submit that claims 12, 15-19, 40, 51, and 117 are definite.

With respect to claims 20 and 97, Applicants amended the claims to include “bits” according to the Examiner’s suggestion. This amendment is not a narrowing amendment. Applicants respectfully submit that claims 20 and 97 are now definite.

Claims 1-11, 13-15, 20, 21-34, 40, 79-89, 91-93, and 97 are rejected under 35 U.S.C. § 112, second paragraph, as being indefinite for failing to particularly point and distinctly claim the subject matter which Applicant regards as the invention.

With respect to claims 1, 6, 21, 29, and 32, the Examiner alleges that the limitation “and that generates a first address from said first program read request” is unclear. Applicants respectfully disagree. The cited portion of claim 1 recites a first line cache interface **that is associated** with the first CPU, **that receives** a first program read request from the first CPU **and that generates** a first address from said first program read request. Applicants respectfully submit that it is clear from this limitation that the first line cache interface generates the first address and that claims 1, 6, 21, 29, and 32 are definite.

With respect to claim 40, the Examiner alleges that the phrase “replacement based on usage of said pages” is unclear because it is unclear “what kind of usage applicant meant.” Applicants respectfully disagree. For example, “[t]he examiner’s focus during examination of the claims for compliance with the requirement for definiteness of 35 U.S.C. 112, second paragraph, is whether the claim meets the threshold requirements of clarity and precision, not whether more suitable language or modes of expression are available.” MPEP § 2173.02. Further, “[i]f the scope of the subject matter embraced by the claims is clear, and if applicants have not otherwise indicated that they intend the invention to be of a scope different from that defined in the claims, then the claims comply with 35 U.S.C. 112, second paragraph.” MPEP § 2173.04.

Applicants respectfully note that “use” and “usage” of cache pages are discussed throughout the claims and specification. For example, “[t]he least used page select device 330 tracks usage of pages stored in cache.” (See Paragraph [0066]). As such, Applicants respectfully submit that claim 40 is definite.

With respect to claim 79, the Examiner alleges that it is unclear whether the line cache receives the first address as well as a second translated address based on the first address. Applicants amended the claim to clarify this limitation. This amendment is not a narrowing amendment. Applicants respectfully submit that claims 79-89, 91-93, and 97 are now definite.

REJECTION UNDER 35 U.S.C. § 103

Claims 1-5, 11, 13-17, 20, 44-48, 50, 52-56, 67-68, 70-71, 73, 79-83, 89, 91-94, and 97 are rejected under 35 U.S.C. § 103(a) as being unpatentable in view of Zaidi et al (U.S. Pat. No. 6,601,126) and Jim Handy (The Cache Memory Book, second edition, published 1998) and Taylor et al. (U.S. Pat. No. 5,699,551). Claims 6-10, 21-28, 30-33, 49, 59-63, 65-66, 84-88, 98-105, 107-110 are rejected under 35 U.S.C. § 103(a), as being unpatentable in view of Zaidi and Jim Handy as applied to claims 1-5, 11, 13-14, 16-17, 44-48, and 79-83 above, and further in view of Barroso et al (U.S. Pat. No. 6,725,334). These rejections are respectfully traversed.

With respect to claim 1, Zaidi either singly or in combination with Handy, Barroso, and Taylor, fails to show, teach, or suggest a line cache that receives a second address that is based on the first address and includes a memory select portion, and a switch

that receives the second address and that selectively connects said line cache to one of said first and second memory interfaces based on the memory select portion.

It is a longstanding rule that to establish a prima facie case of obviousness of a claimed invention, all of the claim limitations must be taught or suggested by the prior art. *In re Royka*, 180 USPQ 143 (CCPA 1974), see MPEP §2143.03. Here, the Examiner fails to provide **any** reference to support a finding that a second address that is based on the first address and includes a memory select portion, and selectively connecting based on the memory select portion are obvious. Furthermore, when evaluating claims for obviousness under 35 U.S.C. §103, all of the limitations must be considered and given weight. *Ex parte Grasselli*, 231 USPQ 393 (Bd. App. 1983), MPEP § 2144.03. Here, it is clear that the Examiner has given little or no consideration of the limitations **and failed to give the limitations any weight**.

An exemplary embodiment of Applicants' invention shown in FIG. 3 illustrates a switch 64 that selectively connects a line cache 58 to one of a buffer memory and a flash memory based on a memory select portion 134 (or 134'). A translated address 130 (or 130') includes the memory select portion 134 (134'). For example, "[w]hen the buffer memory 79 is selected, the translated address 130 includes a memory select portion 134, which selects the buffer memory 79 as the target memory." (Paragraph [0046]). Similarly, "[w]hen the flash memory 86 is selected, the translated address 130' includes a memory select portion 134', which selects the flash memory 86 as the target memory." (Paragraph [0047]). The switch 64 selectively connects the line cache 58 to the appropriate memory device (i.e. outputs the translated address) based on the memory select portion 134 or 134'.

The Examiner acknowledges that Zaidi fails to disclose this limitation, and instead relies on Taylor. In particular, the Examiner cites Column 1, Lines 26-40 of Taylor. The cited portion of Taylor states:

Most current computer systems use virtual memory to provide protection, large address spaces and convenient allocation of physical memory. In virtual memory systems, addresses must be translated at some point between the processor and main memory. While most aspects of cache design are unrelated to virtual memory, whether translation occurs before or after the cache access is a crucial issue and one which is quite visible to the memory management software. If fast cache access were the only consideration, then a virtual cache would be attractive because it could be accessed without waiting for the translation.

Applicants respectfully note that the cited portion merely discusses address translation and is absent of any teaching or suggestion of a second address that includes a memory select portion. As such, the cited portion appears to be absent of any teaching or suggestion of a switch that receives the second address and that selectively connects said line cache to one of said first and second memory interfaces based on the memory select portion.

Here again, **all** of the claim limitations must be taught or suggested by the prior art, and the Examiner fails to provide any prior art reference that discloses these limitations. Applicants respectfully submit that claim 1, as well as its dependent claims, should be allowable for at least the above reasons. Claims 21, 44, 51, 79, and 98, as well as their corresponding dependent claims, should be allowable for at least similar reasons.

With respect to claim 16, the cited art (in particular, Jim Handy) fails to show, teach, or suggest a least used page device that identifies a least used page and a second least used page.

As described in an exemplary embodiment in Paragraph [0072], “the least used page algorithm may be used to identify a least used page and a second least used page when a miss occurs...If the least used page will probably be used by the CPU soon, the second least used page can be replaced instead of the least used page.” In other words, the algorithm returns more than one least used page (i.e. first and second least used pages).

As best understood by Applicants, Jim Handy fails to disclose this limitation. Instead, Jim Handy discloses returning only one value. For example, the Examiner cites page 57, paragraph 2. The cited portion of Jim Handy states:

Some cache controllers watch over the accesses into the cache, and categorize the order in which each Way was accessed, making a note of which Way's line was accessed least recently. This is called a least recently used (LRU) algorithm.”

Applicants respectfully note that the cited portion appears to be absent of any teaching or suggestion of identifying first **and** second least used pages. Instead, the Examiner alleges that this limitation is inherent.

The fact that a certain characteristic **may occur or be present** in the prior art reference is not sufficient to establish inherency of that characteristic. *In re Rijckaert*, 28 USPQ2d 1955, 1957 (Fed. Cir. 1993) (emphasis added). The Federal Circuit has clearly stated that:

To establish inherency, the extrinsic evidence ‘must make clear that the missing descriptive matter is **necessarily** present in the thing described in the reference, and that it would be so recognized by persons of ordinary skill. Inherency, however, may not be established by probabilities or possibilities.’

In re Robertson, 49 USPPQ2d 1949, 1950-1951 (Fed. Cir. 1999) (emphasis added).

“In relying upon the theory of inherency, the examiner must provide a basis in fact and/or technical reasoning to reasonably support the determination that the allegedly inherent characteristic necessarily flows from the teachings of the applied prior art.” *Ex Parte Levy*, 17 USPQ2d 1461 (Bd. Pat. App. & Inter. 1990) (emphasis original). Therefore, identifying first and second least used pages must **necessarily** flow from the teachings of the Jim Handy reference. Applicants respectfully assert that this is not the case here and that the Examiner has failed to properly support his rejection an inherency argument.

Applicants respectfully submit that claim 16, as well as its dependent claims, should be allowable for at least the above reasons. Claims 55, 67, 70, and 94, as well as their corresponding dependent claims, should be allowable for at least similar reasons.

Claims 19, 58, and 96 are rejected under 35 U.S.C. § 103(a), as being unpatentable in view of Zaidi, Jim Handy and Barroso as applied to claims 1-10 above, and further in view of Goodsell (U.S. Pat. No. 7,047,387) and Kirovski et al. (Application-Driven Synthesis of Memory-Intensive System-on-Chip).

With respect to claim 19, the alleged combination fails to show, teach, or suggest a CPU that executes an application, wherein said line cache has a line width and number of pages that are based on said application.

The Examiner alleges that Kirovski discloses this limitation at page 1316, paragraphs 2 and 5. Applicants respectfully note that the cited portion of Kirovski appears to be absent of any teaching or suggestion of a line cache that has a line width and number of pages based on said application. Here again, Applicants respectfully

note that to establish a prima facie case of obviousness of a claimed invention, **all** of the claim limitations must be taught or suggested by the prior art.

For example, the cited portion states that “application-specific integrated circuit (ASIC) floorplans reveal that most of the area in modern SOC’s is dedicated to the processor core and associated caches.” Further, the “Stanford on-line Cache Design Tool...is used to estimate cache access latency and area based on properties such as cache total size, line size, feature size, replacement policy, and set associativity.” Applicants respectfully note that the cited portions do not disclose the limitations of claim 19. Claim 19 requires, specifically, that the line cache has a line width and number of pages based on said application. In contrast, the cited portion is generally directed to cache design and estimating cache access latency. The cited portions do not appear to have any relation to basing line width and number of pages on an application executed by a CPU.

Applicants respectfully submit that claim 19 should be allowable for at least the above reasons. Claims 58 and 96 should be allowable for at least similar reasons.

Claims 18, 34, 35-38, 57, 69, 72, 95, 111 and 112-115 are rejected under 35 U.S.C. § 103(a) as being unpatentable over Zaidi, Jim Handy and Barroso as applied to claims 1-20 above and further in view of Smith et al. (U.S. Pat. No. 5,594,886). This rejection is respectfully traversed.

With respect to claim 35, the alleged combination fails to show, teach, or suggest wherein said line cache module includes a least used page device that selects first and second pages to be overwritten and wherein said line cache module selects one of said first and second pages based on said internal state of the first CPU.

The Examiner relies on Column 1, Lines 59-67 of Smith to disclose this limitation.

The cited portion states:

Some cache controllers use a pseudo least recently used algorithm (pseudo-LRU) to determine the most likely least recently used cache line. Such a cache controller would **determine the last used cache line and select the cache line next to the last used cache line as the likely least used cache line and evict that cache line**. While this algorithm is attractive in its simplicity, it suffers from the fact that the cache line immediately next to the last used cache line is likely not the least used cache line because of the way data is stored in a cache. At best, this algorithm can only ensure the last used cache line is not evicted from the cache. (Emphasis added)

In other words, the cited portion discloses determining the last used cache line and selecting the next to last used cache line as the least used cache line. This is not analogous to selecting first and second pages to be overwritten and selecting one of said first and second pages based on said internal state of the first CPU. The Examiner fails to provide any prior art reference that discloses this limitation.

Applicants respectfully submit that claims 18, 34, 35-38, 57, 69, 72, 95, 111 and 112-115 should be allowable for at least the above reasons.

Claims 12, 29, 39-43, 51, 64-66, 74-78, 90, 106-108, and 116-120 are rejected under 35 U.S.C. § 103(a) as being unpatentable over Zaidi, Jim Handy, Barroso, and Smith as applied to claims 1-11 and 13-20 above and further in view of Dottling et al. (U.S. Pat. No. 6,014,756). This rejection is respectfully traversed.

With respect to claim 12, the alleged combination fails to show, teach, or suggest said line cache memory includes multiple pages and wherein said line cache module allows one page to be accessed by one of the first CPU and said second CPU while the other of the first CPU and said second CPU is waiting for data retrieval in another page.

The Examiner alleges that Column 1, Lines 38-56 of Dottling discloses this limitation. Applicants respectfully disagree. Here again, **all** of the claim limitations must be taught or suggested by the prior art. For example, the cited portion of Dottling states:

Further EP-A-0 637 799 is directed to a high performance shared cache for multiprocessor systems, which supports the multiprocessor system insofar as allowing maximum parallelism in accessing the cache by the processing units, **servicing one processor request in each machine cycle**, reducing system response time, and increasing system throughput. The shared cache disclosed therein, uses the additional performance optimization techniques of pipelining cache operations, i.e. loads and stores, and burst-mode data accesses. By including built-in pipeline stages, **that cache is enabled to service one request every machine cycle from any processing element**. This contributes to reduction in the system response time as well as the throughput. Particularly one portion of the data is held by means of a logic circuitry of the cache, while another portion, corresponding to the system bus width, gets transferred to the requesting element, e.g. a processor or a storage, in one cycle. The held portion of the data can then be transferred in the following machine cycle. (Emphasis added)

Applicants respectfully note that the above portion appears to be absent of any teaching or suggestion of allowing one page to be accessed by one of the first CPU and said second CPU while the other of the first CPU and said second CPU is waiting for data retrieval in another page. Simply disclosing shared cache in a multiprocessor system is not analogous to allowing one page to be accessed by one CPU while the other CPU waits for data retrieval in another page.

Applicants respectfully submit that claims 12, 29, 39-43, 51, 64-66, 74-78, 90, 106-108, and 116-120 should be allowable for at least the above reasons.

CONCLUSION

It is believed that all of the stated grounds of rejection have been properly traversed, accommodated, or rendered moot. Applicant therefore respectfully requests that the Examiner reconsider and withdraw all presently outstanding rejections. It is believed that a full and complete response has been made to the outstanding Office Action and the present application is in condition for allowance. Thus, prompt and favorable consideration of this amendment is respectfully requested. If the Examiner believes that personal communication will expedite prosecution of this application, the Examiner is invited to telephone the undersigned at (248) 641-1600.

Respectfully submitted,

Dated: September 29, 2006

By: 

Damian M. Aquino
Reg. No. 54,964

HARNESS, DICKEY & PIERCE, P.L.C.
P.O. Box 828
Bloomfield Hills, Michigan 48303
(248) 641-1600

MDW/DMA/mp